

We claim:

- Sub*
- AI*
1. A system for storing data, the system having one or more storage devices, the system comprising:
10 a first cache level for caching data from a sender into a first random-access structure;
 a second cache level for caching data from the first cache level into a log structure; and
 a storage level for storing data from CL into a second random-access structure, wherein
CL is the first cache level or the second cache level.
- 15 2. The system according to claim 1 wherein the log structure includes one or more segments.
3. The system according to claim 2 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.
- 20 4. The system according to claim 2 wherein the log structure further includes a segment summary, the segment summary tracking information regarding the configuration of at least one segment in the log structure.
- 25 5. The system according to claim 2 wherein the log structure further includes a bit map, the bit map tracking information regarding live user data stored in at least one segment in the log structure.
- 30 6. The system according to claim 2 wherein the log structure further includes one or more logical-to-physical maps, the logical-to-physical maps tracking information regarding the location of user data stored in at least one segment in the log structure.

7. The system according to claim 2 wherein the log structure further includes one or more map elements, the map elements tracking information regarding the location of contiguous user data stored in at least one segment in the log structure.

8. The system according to claim 2 wherein the log structure further includes one or more 10 segment elements.

9. The system according to claim 8 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

10. The system according to claim 9 wherein the segment database is stored in at least one segment element.

11. The system according to claim 2 wherein the log structure further includes one or more 20 segment blocks.

12. The system according to claim 11 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

25 13. The system according to claim 12 wherein the segment database is stored in at least one segment block.

14. The system according to claim 1 wherein the first cache level is a memory-based cache.

30 15. The system according to claim 1 wherein the first cache level is a write-back cache.

16. The system according to claim 1 wherein the second cache level is a disk-based cache.
17. The system according to claim 1 wherein the second cache level is a write-back cache.
18. The system according to claim 1 wherein the second cache level further caches in the log structure parity data for the data cached in the log structure.
19. The system according to claim 1 wherein the storage level further stores in the second random-access structure parity data for the data stored in the second random-access structure.
20. The system according to claim 1 wherein the log structure is configured as a Redundant Array of Independent Disks structure.
21. The system according to claim 1 wherein the second random-access structure is configured as a Redundant Array of Independent Disks structure.
22. The system according to claim 1 wherein the storage capacity of the first cache level is smaller than the storage capacity of the second cache level.
23. The system according to claim 1 wherein the storage capacity of the second cache level is at least 10% of the storage capacity of the one or more storage devices.
24. The system according to claim 1 wherein the one or more storage devices are disks.
25. The system according to claim 1 wherein the one or more storage devices are a disk array.

26. A method for storing data on a system, the system having one or more storage devices, the method comprising:

caching data from a sender into a first random-access structure located in a first cache level;

caching data from the first cache level into a log structure located in a second cache level;

10 and

storing data from CL into a second random-access structure located in a storage level, wherein CL is the first cache level or the second cache level.

15 27. The method according to claim 26 wherein the log structure includes one or more segments.

20 28. The method according to claim 27 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

25 29. The method according to claim 27 wherein the log structure further includes a segment summary, the segment summary tracking information regarding the configuration of at least one segment in the log structure.

30 30. The method according to claim 27 wherein the log structure further includes a bit map, the bit map tracking information regarding live user data stored in at least one segment in the log structure.

35 31. The method according to claim 27 wherein the log structure further includes one or more logical-to-physical maps, the logical-to-physical maps tracking information regarding the location of user data stored in at least one segment in the log structure.

32. The method according to claim 27 wherein the log structure further includes one or more map elements, the map elements tracking information regarding the location of contiguous user data stored in at least one segment in the log structure.

33. The method according to claim 27 wherein the log structure further includes one or more segment elements.

34. The method according to claim 33 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

35. The method according to claim 34 wherein the segment database is stored in at least one segment element.

36. The method according to claim 27 wherein the log structure further includes one or more segment blocks.

37. The method according to claim 36 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

38. The method according to claim 37 wherein the segment database is stored in at least one segment block.

39. The method according to claim 26 wherein the first cache level is a memory-based cache.

40. The method according to claim 26 wherein the first cache level is a write-back cache.

41. The method according to claim 26 wherein the second cache level is a disk-based cache.

42. The method according to claim 26 wherein the second cache level is a write-back cache.

10 43. The method according to claim 26 wherein the second cache level further caches in the log structure parity data for the data cached in the log structure.

44. The method according to claim 26 wherein the storage level further stores in the second random-access structure parity data for the data stored in the second random-access structure.

15 45. The method according to claim 26 wherein the log structure is configured as a Redundant Array of Independent Disks structure.

46. The method according to claim 26 wherein the second random-access structure is configured as a Redundant Array of Independent Disks structure.

20 47. The method according to claim 26 wherein the storage capacity of the first cache level is smaller than the storage capacity of the second cache level.

48. The method according to claim 26 wherein the storage capacity of the second cache level 25 is at least 10% of the storage capacity of the one or more storage devices.

49. The method according to claim 26 wherein the one or more storage devices are disks.

50. The method according to claim 26 wherein the one or more storage devices are a disk 30 array.

51. A system for storing data, the system having one or more storage devices, the system comprising:

means for caching data from a sender into a first random-access structure located in a first cache level;

means for caching data from the first cache level into a log structure located in a second cache level; and

means for storing data from CL into a second random-access structure located in a storage level, wherein CL is the first cache level or the second cache level.

52. The system according to claim 51 wherein the log structure includes one or more segments.

53. The system according to claim 52 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

54. The system according to claim 52 wherein the log structure further includes a segment summary, the segment summary tracking information regarding the configuration of at least one segment in the log structure.

- 25 55. The system according to claim 52 wherein the log structure further includes a bit map, the
bit map tracking information regarding live user data stored in at least one segment in the log
structure.

56. The system according to claim 52 wherein the log structure further includes one or more
30 logical-to-physical maps, the logical-to-physical maps tracking information regarding the
location of user data stored in at least one segment in the log structure.

57. The system according to claim 52 wherein the log structure further includes one or more map elements, the map elements tracking information regarding the location of contiguous user data stored in at least one segment in the log structure.

58. The system according to claim 52 wherein the log structure further includes one or more segment elements.

59. The system according to claim 58 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

60. The system according to claim 59 wherein the segment database is stored in at least one segment element.

61. The system according to claim 52 wherein the log structure further includes one or more segment blocks.

62. The system according to claim 61 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

63. The system according to claim 62 wherein the segment database is stored in at least one segment block.

64. The system according to claim 51 wherein the first cache level is a memory-based cache.

65. The system according to claim 51 wherein the first cache level is a write-back cache.

66. The system according to claim 51 wherein the second cache level is a disk-based cache.

67. The system according to claim 51 wherein the second cache level is a write-back cache.

68. The system according to claim 51 wherein the second cache level further caches in the log
10 structure parity data for the data cached in the log structure.

69. The system according to claim 51 wherein the storage level further stores in the second
random-access structure parity data for the data stored in the second random-access structure.

70. The system according to claim 51 wherein the log structure is configured as a Redundant
Array of Independent Disks structure.

71. The system according to claim 51 wherein the second random-access structure is
configured as a Redundant Array of Independent Disks structure.

72. The system according to claim 51 wherein the storage capacity of the first cache level is
smaller than the storage capacity of the second cache level.

73. The system according to claim 51 wherein the storage capacity of the second cache level
25 is at least 10% of the storage capacity of the one or more storage devices.

74. The system according to claim 51 wherein the one or more storage devices are disks.

75. The system according to claim 51 wherein the one or more storage devices are a disk
30 array.

76. A computer program product for storing data on a system, the system having one or more storage devices, the computer program product comprising computer usable medium, the computer usable medium having computer readable program code thereon, the computer readable program code comprising:

10 program code for caching data from a sender into a first random-access structure located in a first cache level;

15 program code for caching data from the first cache level into a log structure located in a second cache level; and

20 program code for storing data from CL into a second random-access structure located in a storage level, wherein CL is the first cache level or the second cache level.

25 77. The computer program product according to claim 76 wherein the log structure includes one or more segments.

30 78. The computer program product according to claim 77 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

79. The computer program product according to claim 77 wherein the log structure further includes a segment summary, the segment summary tracking information regarding the configuration of at least one segment in the log structure.

80. The computer program product according to claim 77 wherein the log structure further includes a bit map, the bit map tracking information regarding live user data stored in at least one segment in the log structure.

81. The computer program product according to claim 77 wherein the log structure further includes one or more logical-to-physical maps, the logical-to-physical maps tracking information regarding the location of user data stored in at least one segment in the log structure.

82. The computer program product according to claim 77 wherein the log structure further includes one or more map elements, the map elements tracking information regarding the location of contiguous user data stored in at least one segment in the log structure.

83. The computer program product according to claim 77 wherein the log structure further includes one or more segment elements.

84. The computer program product according to claim 83 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

85. The computer program product according to claim 84 wherein the segment database is stored in at least one segment element.

86. The computer program product according to claim 77 wherein the log structure further includes one or more segment blocks.

87. The computer program product according to claim 86 wherein the log structure further includes a segment database, the segment database tracking information regarding the configuration of user data stored in at least one segment in the log structure.

88. The computer program product according to claim 87 wherein the segment database is stored in at least one segment block.

89. The computer program product according to claim 76 wherein the first cache level is a memory-based cache.

90. The computer program product according to claim 76 wherein the first cache level is a write-back cache.

10 91. The computer program product according to claim 76 wherein the second cache level is a disk-based cache.

92. The computer program product according to claim 76 wherein the second cache level is a write-back cache.

93. The computer program product according to claim 76 wherein the second cache level further caches in the log structure parity data for the data cached in the log structure.

20 94. The computer program product according to claim 76 wherein the storage level further stores in the second random-access structure parity data for the data stored in the second random-access structure.

25 95. The computer program product according to claim 76 wherein the log structure is configured as a Redundant Array of Independent Disks structure.

96. The computer program product according to claim 76 wherein the second random-access structure is configured as a Redundant Array of Independent Disks structure.

30 97. The computer program product according to claim 76 wherein the storage capacity of the first cache level is smaller than the storage capacity of the second cache level.

98. The computer program product according to claim 76 wherein the storage capacity of the second cache level is at least 10% of the storage capacity of the one or more storage devices.

99. The computer program product according to claim 76 wherein the one or more storage devices are disks.

10

100. The computer program product according to claim 76 wherein the one or more storage devices are a disk array.

15 A system for storing data, the system having one or more storage devices, the system comprising:

a cache level for caching data from a sender, the cached data written to the cache level in a sequential manner; and

a storage level for storing data from CL, wherein CL is the sender or the cache level, the stored data written to the storage level in a random manner.

20

102. The system according to claim 101 wherein the cache level is a disk-based cache.

103. The system according to claim 101 wherein the cache level is a write-back cache.

25

104. The system according to claim 101 wherein the cache level further caches parity data for the data cached in the cache level.

105. The system according to claim 101 wherein the storage level further stores parity data for the data stored in the storage level.

30

106. The system according to claim 101 wherein the cache level is configured as a Redundant Array of Independent Disks structure.

107. The system according to claim 101 wherein the storage level is configured as a Redundant Array of Independent Disks structure.

108. The system according to claim 101 wherein the storage capacity of the cache level is at least 10% of the storage capacity of the one or more storage devices.

109. The system according to claim 101 wherein the one or more storage devices are disks.

110. The system according to claim 101 wherein the one or more storage devices are a disk array.

111. A method for storing data on a system, the system having one or more storage devices, the method comprising:

20 caching data from a sender to a cache level, the cached data written to the cache level in a sequential manner; and

25 storing data from CL to a storage level, wherein CL is the sender or the cache level, the stored data written to the storage level in a random manner.

112. The method according to claim 111 wherein the cache level is a disk-based cache.

113. The method according to claim 111 wherein the cache level is a write-back cache.

30 114. The method according to claim 111 wherein the cache level further caches parity data for the data cached in the cache level.

115. The method according to claim 111 wherein the storage level further stores parity data for the data stored in the storage level.

116. The method according to claim 111 wherein the cache level is configured as a Redundant Array of Independent Disks structure.

10 117. The method according to claim 111 wherein the storage level is configured as a Redundant Array of Independent Disks structure.

118. The method according to claim 111 wherein the storage capacity of the cache level is at least 10% of the storage capacity of the one or more storage devices.

119. The method according to claim 111 wherein the one or more storage devices are disks.

20 120. The method according to claim 111 wherein the one or more storage devices are a disk array.

25 121. A computer program product for storing data on a system, the system having one or more storage devices, the computer program product comprising a computer usable medium, the computer usable medium having computer readable code thereon, the computer readable code comprising:

program code for caching data from a sender, the cached data written to the cache level in a sequential manner; and

program code for storing data from CL, wherein CL is the sender or the cache level, the stored data written to the storage level in a random manner.

122. The computer program product according to claim 121 wherein the cache level is a disk-based cache.
123. The computer program product according to claim 121 wherein the cache level is a write-back cache.
- 10 124. The computer program product according to claim 121 wherein the cache level further caches parity data for the data cached in the cache level.
- 15 125. The computer program product according to claim 121 wherein the storage level further stores parity data for the data stored in the storage level.
- 20 126. The computer program product according to claim 121 wherein the cache level is configured as a Redundant Array of Independent Disks structure.
- 25 127. The computer program product according to claim 121 wherein the storage level is configured as a Redundant Array of Independent Disks structure.
128. The computer program product according to claim 121 wherein the storage capacity of the cache level is at least 10% of the storage capacity of the one or more storage devices.
- 30 129. The computer program product according to claim 121 wherein the one or more storage devices are disks.
130. The computer program product according to claim 121 wherein the one or more storage devices are a disk array.